





Virtualization and Disaster Recovery: Implementing and Automating Disaster and Data Recovery for z/VM

Daniel P. Martin
Rocket Software, Inc.

<dan.martin@rocketsoftware.com>


Thursday, 14 August 2008
Session Number 9157





*When going to hide, know how to get there...
...and how to get back...
...and eat first.*

-- Stephen Sondheim, "Into the Woods"



SHARE
Technology • Connections • Results


An Agenda

- Introductions: *Who is this person, and why is he here?*
- Terminology: *A very brief review of virtualization.*
- More terminology: *Disaster Recovery versus Data Recovery.*
- Tools: *Machinery is aggressive.*
- Automation: *I know what I said. Do what I meant.*
- Specifics: *Eventually you have to stop planning and start "doing".*
- An example case: *"Meanwhile, back in the lab..."*
- For further reading...

Introductions


- **About Rocket Software**

- Rocket Software was founded in 1990 with an initial focus on developing and delivering software solutions to IBM mainframe customers deploying IBM DB2. Within the first five years, Rocket licensed all of its intellectual property to IBM, and all of Rocket's products were rebranded and licensed to enterprises through IBM's worldwide sales channel.
- Since those early years, Rocket has expanded its software OEM business to cover a wide spectrum of enterprise infrastructure products in the areas of [Business Intelligence](#), [Storage](#), [Networks](#), [Terminal Emulation](#), [Integration](#), [Security](#), and [Database](#). And in addition to IBM, Rocket has licensed its products to a number of major [OEMs](#) that include Microsoft, RSA, EMC, Nortel, Motorola, and HP.
- In parallel with a significant R&D investment in developing software products, Rocket Software began a series of strategic acquisitions in 2000. As a result, Rocket now has well over 100 software products that are licensed both through OEM partners and resellers, and via direct sales through Rocket's acquired branded sales channels, such as [Mainstar](#), [Servergraph](#), [CorVu](#), [BlueZone](#), and [Seagull](#).
- Rocket values the technical skills and deep domain expertise of the employees at the companies we acquire. We endeavor to keep the teams intact and entrust them to manage and expand the product lines that they initially created. This philosophy has allowed Rocket to grow to more than 700 employees around the world, including Eastern and Western Europe, Australia, Asia and North America.



The details of my life are quite inconsequential...


-- Dr. Evil



SHARE
Technology • Connections • Results

Introductions

- About your speaker
 - Life as a customer...
 - z/VM System Programmer starting some time around VM/370 R6 PLC 3
 - *i.e., "For quite a while..."*
 - University of Arkansas
 - *"VM Workshop" speaker and tools tape contributor; IBMVM (nee VMESA-L) List Moderator*
 - *"Sir Dan, Keeper of the List"*
 - *Unix/Linux sysadmin; Open systems / IT security Manager*
 - *See above, re: "For quite a while..."*
 - Multiple platforms
 - *z/OS ancestors, MUSIC, UTS, OS/2, Unix, Solaris, Linux, Beowulf Clustering*
 - *Assorted strange and fierce creatures: Plan 9, Teradata, DPCX, DPPX, Amoeba*
 - *s/370 et cetera, 3650, 8100; Intel, POWER, SPARC, MIPS, PA-RISC, other...*
 - *i.e., "I've done some weird stuff..."*
 - Life as a developer...
 - Senior Software Developer for Rocket Software, Inc since 2003
 - Lead developer, IBM Backup and Restore Manager for z/VM (5697-J06) and IBM Archive Manager for z/VM (5697-J05)
 - Team Member for development of IBM Tape Manager for z/VM (5697-J08) and IBM Operations Manager for z/VM (5697-J10)
 - Life away from the keyboard...
 - First responder and sworn law enforcement officer
 - *"...but, no, I can't fix that speeding ticket for you."*




*If 'real' is what you can feel, smell, taste and see, then
'real' is simply electrical signals interpreted by your
brain.*

-- Morpheus, "The Matrix"




Terminology

- Virtualization: If you can see it, but it isn't really there...
 - Many possible approaches
 - Full, Hardware-assisted, Partial, Paravirtualization, OS-level, Resource, Application...
 - Modern z/Architecture systems are virtualized on several fronts:
 - LPAR
 - Storage Subsystems
 - Network Interfaces
 - ...and we're not even going to consider talking about the implications of processor microcode.
 - For purposes of the present discussion
 - As implemented by z/VM operating within a z/Architecture hardware environment.
 - "Forty-one-derful years"
 - Mother knows best – <http://www.princeton.edu/~melinda/25paper.pdf> (SHARE 89, Aug '97)



Court disaster long enough, and it will accept your proposal.


-- Mason Cooley



SHARE
Technology • Connections • Results


More Terminology

- Disaster: A calamitous event
 - Data: Pieces of information
 - Recovery: Regaining something lost or taken away
-
- *What on earth is he talking about!?*



With any recovery from morbidity there must go a certain healthy humiliation.


-- G.K. Chesterton



SHARE
Technology • Connections • Results

More Terminology

- In the present context...
- Disaster Recovery
 - Generally: The ability to resume business operations
 - “Name your favorite catastrophic event.”
 - Specifically: Getting your z/VM system back “on air” after a catastrophic interruption
 - See also: Two excellent disaster recovery “case study” sessions: Session 9137, “Using VM for Linux Disaster Recovery Planning” *and* Session 9210, “z/VM and Linux Disaster Recovery – A Customer Experience”.



Experience is the name every one gives to their mistakes.

-- Oscar Wilde



More Terminology


- In the present context...
- Data Recovery
 - Generally: The ability to recover “lost” information
 - “Oops...” and what comes afterward
 - Specifically: Resuming some part of business operations after a lesser catastrophic event
 - Usually: Reduced impact versus full-blown disaster
 - A single file
 - A single minidisk
 - An application



Tools

Machinery is aggressive. The weaver becomes a web, the machinist a machine. If you do not use the tools, they use you. All tools are in one sense edge-tools, and dangerous.

-- Ralph Waldo Emerson



SHARE
Technology • Connections • Results


- Numerous different tools available in the marketplace
 - IBM
 - Other vendors
 - Open source
 - Roll-your-own
- Regardless of choice, it's vitaly important to
 - Know your site policies and requirements
 - Know your systems and applications
 - Know your tools
 - Know your budget
- For the case at hand:
 - 5697-J10 IBM Operations Manager for z/VM 1.3.0 (GA June '08; no PTF)
 - 5697-J06 IBM Backup and Restore Manager for z/VM 1.2.0 (+ PTF UK36984)
 - 5697-J08 IBM Tape Manager for z/VM 1.2.0 (+ PTF UK38722)
 - 5697-J05 IBM Archive Manager for z/VM 1.1.0 (+ PTF UK32746)



Automation

Skilled workers historically have been ambivalent toward automation, knowing that the bodies it would augment or replace were the occasion for both their pain and their power.

-- Shoshana Zuboff, "In the Age of the Smart Machine"




SHARE
Technology • Connections • Results

- Why automate?
 - 24x7x365 nature of the modern enterprise
 - Business operations are “always on”; people are not.
 - Computers aren’t prone to boredom.
 - Improved reliability
 - Because people and people-driven processes are fragile.
 - “Shrink-wrapped” procedures reduce opportunity for error.
 - Entropy always wins...
 - *...but we can take steps to improve the odds.*
 - Improved auditability
 - Because “bean counters” like records...
 - *And because there is business value which can not be ignored in keeping auditors, accountants and their kin happy.*
 - Budget considerations
 - “Unanticipated procedural breakdowns” (i.e., errors) can be expensive
 - Even home-grown tools come at a cost

Simplify, simplify.

-- Henry David Thoreau, "Walden"




SHARE
Technology • Connections • Results

Automation

- “Do what I Said” versus “Do What I Meant”
 - You can not automate what you do not understand.

“Clarity is of no importance because nobody listens and nobody knows what you mean no matter what you mean, nor how clearly you mean what you mean. But if you have vitality enough of knowing enough of what you mean, somebody and sometime and sometimes a great many will have to realize that you know what you mean and so they will agree that you mean what you know, what you know you mean, which is as near as anybody can come to understanding any one.”


-- Gertrude Stein (1874–1946)



Automation – An example

It is unworthy of excellent men to lose hours like slaves in the labour of calculation which could safely be relegated to anyone else if machines were used.

-- Gottfried Wilhelm Leibniz (1646 – 1716)




SHARE
Technology • Connections • Results

- IBM Operations Manager for z/VM – Example Rules for Linux Guest File System Backup


- Quiesce
- Backup
- Restart

```
*
* Step 1 - Scheduled shutdown of Linux guest BKRLNXD
*
DEFSCHD NAME (BKRLNX01), WHEN (THU-13:15), ACTION (BKRLNX01)
DEFACTN NAME (BKRLNX01), COMMAND (CP SIGNAL SHUTDOWN BKRLNXD WITHIN 120), +
ENV (LVM)
*
* Step 2 - Detect BKRLNXD shutdown and start backup job SAMPLNX
*
DEFRULE NAME (BKRLNX02), MATCH (*HCPSIG211%I User BKRLNXD*), ACTION (BKRLNX02)
DEFACTN NAME (BKRLNX02), COMMAND (CP SMSG BKRBKUP SUBMIT SAMPLNX), ENV (LVM)
*
* Step 3 - Define virtual machine group BKRWRKRS; detect "backup
*         complete" message at end of SAMPLNX backup job;
*         restart Linux guest BKRLNXD.
*
DEFGROUP NAME (BKRWRKRS), USER (BKRWRK0*)
DEFRULE NAME (BKRLNX03), MATCH (*BACKUP COMPLETE - BKRLNXD*), +
GROUP (BKRWRKRS), ACTION (BKRLNX03)
DEFACTN NAME (BKRLNX03), COMMAND (CP XAUTOLOG BKRLNXD), ENV (LVM)
```



Suffering isn't ennobling, recovery is.

-- Christiaan N. Barnard




SHARE
Technology • Connections • Results

Backup – z/VM view

- We call it “backup” but it’s ultimately all about recovery.
 - Otherwise, what would be the point?
- Know your systems! Know your applications!
 - *Even if you only use z/VM to host Linux (or other) guest operating systems, you need a sufficient backup of your z/VM system to enable recovery of basic system functionality and enterprise-specific customizations.*
 - IBM Backup and Restore Manager enables resource selection by “INCLUDE / EXCLUDE / SELECT” mechanism for identification of DASD (and SFS) resources.
 - Use CP “SPXTAPE” function for backup of SDF (System Data File) content.
 - If you can not recover an IPL’able system, how will you recover your guest operating systems and applications?

A good decision is based on knowledge and not on numbers.

-- Plato




SHARE
Technology • Connections • Results

Backup – Making choices

- **IBM Backup and Restore Manager for z/VM – A sample set of INCLUDE / EXCLUDE definitions**


FUNCTION	MEDIATYPE	OWNER	VDEV	VOLUME	DEVTYPE	START	END	SIZE	RESERVED
Include	Minidisk	*	= *	*	*	= *	= *	= *	*
Exclude	Minidisk	FDISK	= *	*	*	= *	= *	= *	*
Exclude	Minidisk	\$ALLOC\$	= *	*	*	= *	= *	= *	*
Exclude	Minidisk	VMSESV*	= *	*	*	= *	= *	= *	*
Include	Minidisk	VMSESV*	= 0191	*	*	= *	= *	= *	*
Exclude	Minidisk	*	= *	*	*	= *	= *	> 3300	*
Exclude	Minidisk	*	= *	*	*	= *	= *	= END	*
Include	Minidisk	MAINT	= 012*	*	*	= *	= *	= *	*

FUNCTION	MEDIATYPE	POOLNAME	OWNER	FS
Include	SFS	VMSYSU:	*	SFS
Exclude	SFS	VMSYSU:	DFSMS*	SFS
Include	SFS	VMDEVU:	*	SFS



*Things of this world are in so constant a flux, that
nothing remains long in the same state.*



-- John Locke (1632-1704)




SHARE
Technology • Connections • Results

Backup – “moving target” problem

“Write Pointer” – File System Activity




“Read Pointer” – Backup Activity



Excepting death, there are no great disasters.


-- Chinese proverb.



S H A R E
Technology • Connections • Results

Data Recovery vs. Disaster Recovery

- Context: The Virtualized Perspective
- Native z/VM Resources
 - CMS data (minidisk files, SFS) are susceptible to both full and incremental backup processing.
 - IBM Backup and Restore Manager (and other products) have file system awareness.
 - Usual approach
 - *One (or more) backup job(s) for native resources*
 - *One (or more) backup process for Disaster Recovery purposes*
 - DASD (CMS minidisk or SFS) repository for backup data may be a useful alternative to tape.
 - DDR-format (a/k/a DDRTAPE) backup for DR purposes
 - Cold-site restore enabled via stand-alone DDR
 - *See “z/VM CP Commands and Utility Reference”, the “UTILITY” command*
 - *Remember that DDR does not understand native CMS file system formats. DDR only deals with track- or block-image backups.*



Those who escape death in great disasters are surely
destined for good fortune later.

-- Chinese Proverb



Data Recovery vs. Disaster Recovery

- Context: The Virtualized Perspective
- Guest operating system resources
 - The z/VM hypervisor components (in this case, CP *plus* CMS programming interfaces) are unaware of guest OS file systems.
 - For disaster recovery purposes, backup is usually performed at the *media* level, **not** the *file system* level.
 - IBM Backup and Restore Manager (and other products) require a similar approach for non-native data – track- or block-image backup.
 - Depending on available resources, and on requirements for data recovery versus disaster recovery, IBM Backup and Restore Manager users can use Tape, DASD, SFS or DDRTAPE backup formats.




Tape Management

In early days, I tried not to give librarians any trouble, which was where I made my primary mistake. Librarians like to be given trouble; they exist for it, they are geared to it. For the location of a mislaid volume, an uncatalogued item, your good librarian has a ferret's nose. Give her a scent and she jumps the leash, her eye bright with battle.
-- Catherine Bowen




- Common choice of storage media for backup and recovery
 - High data density
 - Up to 300GB per cartridge (uncompressed) with 3592 media
 - Relatively low cost per gigabyte
 - Necessitates a “librarian” role
- Highly portable
 - Logical portability: “Long-haul channel” connections to remote robotics
 - Physical portability
 - Remote vaulting
 - Recovery site



Logical consequences are the scarecrows of fools and the beacons of wise men.


-- Thomas Henry Huxley



SHARE
Technology • Connections • Results

Tape Management

- Be aware of unintended consequences
 - You can not automate for every eventuality; have a plan for dealing with...
 - Media destruction
 - Physical damage
 - Loss of logical integrity
 - Hazards of physical portability
 - “Oops, I lost my data” makes about as much sense as “Oops, I lost my hair.”
 - *You didn’t “lose” your data. Your media wound up in an unintended location.*
 - Data longevity
 - Expiration from the backup (or archive) catalog does not mean that the bits vanish from the media.
 - Encryption
 - Is generally a good thing, but nothing comes for free
 - Is not a panacea
 - Is an ongoing “arms race”
 - Imposes additional disaster recovery considerations




The sage does not hoard. Having bestowed all he has on others, he has yet more; Having given all he has to others, he is richer still.

-- Lao-Tzu




Archive Management

- Functionally similar to backup, except...
 - Archive transactions tend to have an ad-hoc nature
 - Different approach to content management and retention
 - May not be readily susceptible to automation
 - And will likely require different automation tactics
- Yet another “Know Your System. Know Your Enterprise.” point:
 - If the data had no value, why was the archive created in the first place?
 - Don’t forget about archival data when thinking about disaster recovery.
 - Off-site
 - Twinned media
 - Vaulting
 - Even if your tape vault is “nuke proof”...
 - ...can you recover your backup, archive and tape management catalogs?



*We're all of us guinea pigs in the laboratory of God.
Humanity is just a work in progress.*


-- Tennessee Williams, "Camino Real"



SHARE
Technology • Connections • Results


A Laboratory Scenario (1)

- Automation via IBM Operations Manager for z/VM
 - Scheduling
 - Event monitoring
 - Linux guest management
 - Console activity recording
- Weekly full system backup (“SAMPFULL”)
 - Backup to ATL-resident 3590 tape
 - 30-day retention
 - Tape Management functions handled by IBM Tape Manager for z/VM



*Finally, we come to my number two man. His name?
Number Two.*


-- Dr. Evil



SHARE
Technology • Connections • Results


A Laboratory Scenario (2)

- Weekly disaster recovery backup (“DDRSAMP”)
 - Backup to ATL-resident 3590 (manual off-siting)
 - 30-day catalog retention
 - DDR-format backup via DDRTAPE
- Daily incremental system backup (“SAMPINCR”)
 - Backup to DASD (minidisk and/or SFS destinations)
 - CMS minidisk and SFS files only
- Linux guest file system backup (“SAMPLNX”)
 - Interaction among Operations Manager, Backup and Restore Manager, and Linux guest operating system
- Sample materials are distributed as part of IBM Backup and Restore Manager



Reading is equivalent to thinking with someone else's head instead of with one's own.


-- Arthur Schopenhauer



SHARE
Technology • Connections • Results


Further Reading and Resources

- IBM's z/VM Web Information
 - z/VM Storage Management Tools
 - <http://www.vm.ibm.com/storman/>
- Session 9137, Using VM for Linux Disaster Recovery Planning (Rick Barlow / Nationwide Insurance)
- Session 9210, z/VM and Linux Disaster Recovery – A Customer Experience (Lee Stewart / Sirius Computer Solutions)
- The z/VM Community Discussion – ibmvm@listserv.uark.edu
 - <http://listserv.uark.edu/archives/ibmvm.html>
- And, of course... the entire z/VM Publication Library
 - <http://www.vm.ibm.com/library/>



*Knowing trees, I understand the meaning of patience.
Knowing grass, I can appreciate persistence.*

-- Hal Borland



SHARE
Technology • Connections • Results

And, finally...

THANK YOU!!!

DANIEL P. MARTIN – SENIOR SOFTWARE DEVELOPER

ROCKET SOFTWARE, INC.

1400 NE McCLAIN RD, SUITE 8

BENTONVILLE, AR 72712

VOICE: (479) 271-5280

FAX: (617) 630-7129

DAN.MARTIN@ROCKETSOFTWARE.COM